



US008446818B2

(12) **United States Patent**
Moisiadis et al.

(10) **Patent No.:** **US 8,446,818 B2**
(45) **Date of Patent:** **May 21, 2013**

(54) **ROUTED SPLIT MULTI-LINK TRUNKING
RESILIENCY FOR WIRELESS LOCAL AREA
NETWORK SPLIT-PLANE ENVIRONMENTS**

(75) Inventors: **Alexandros Moisiadis**, Stow, MA (US);
Mohnish Anumala, Littleton, MA (US);
Debin Zhang, Littleton, MA (US);
Lawrence James DiBurro, Haverhill,
MA (US); **Shu Ching Shieh**, Nashua,
NH (US)

(73) Assignee: **Avaya Inc.**, Basking Ridge, NJ (US)

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 397 days.

(21) Appl. No.: **12/917,014**

(22) Filed: **Nov. 1, 2010**

(65) **Prior Publication Data**

US 2012/0106320 A1 May 3, 2012

(51) **Int. Cl.**
H04J 1/16 (2006.01)

(52) **U.S. Cl.**
USPC 370/218; 370/217; 370/219; 370/220;
370/242

(58) **Field of Classification Search** 370/217-220,
370/242

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,463,475	B1 *	10/2002	Calhoun	709/227
7,120,151	B1 *	10/2006	Ginjpalli et al.	370/395.1
2002/0046348	A1 *	4/2002	Brustoloni	713/201
2003/0223358	A1 *	12/2003	Rigby et al.	370/218
2008/0130509	A1 *	6/2008	Pampati et al.	370/241.1
2012/0110393	A1 *	5/2012	Shieh et al.	714/48
2012/0230183	A1 *	9/2012	Allan et al.	370/218

* cited by examiner

Primary Examiner — Hassan Kizou


Assistant Examiner — Ashil Farahmand

(74) *Attorney, Agent, or Firm* — Anderson Gorecki &
Rouille LLP

(57) **ABSTRACT**

A method, apparatus and computer program product for performing Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments is presented. A first network switch learns first data associated with a second network switch and the first network switch terminates an Access Tunnel (AT). The said second network switch learns second data associated with said first network switch. The first and second data are synchronized between the first network switch and the second network switch. The first network switch and the second network load sharing tunnel data packets. The second network switch forwards tunnel control packets received by the second network switch to the first network. A failure relating to said first network switch is detected and a new AT is established with the second network switch.

11 Claims, 6 Drawing Sheets






FIGURE 1




Figure 2







FIGURE 3

**FIGURE 4A**

**FIGURE 4B**

200**Figure 5**

ROUTED SPLIT MULTI-LINK TRUNKING RESILIENCY FOR WIRELESS LOCAL AREA NETWORK SPLIT-PLANE ENVIRONMENTS

BACKGROUND

Data communication networks may include various computers, servers, nodes, routers, switches, hubs, proxies, and other devices coupled to and configured to pass data to one another. These devices are referred to herein as “network elements,” and may provide a variety of network resources on a network. Data is communicated through data communication networks by passing protocol data units (such as packets, cells, frames, or segments) between the network elements over communication links on the network. A particular protocol data unit may be handled by multiple network elements and cross multiple communication links as it travels between its source and its destination over the network. Hosts such as computers, telephones, cellular telephones, Personal Digital Assistants, and other types of consumer electronics connect to and transmit/receive data over the communication network and, hence, are users of the communication services offered by the communication network.

Network elements (e.g. Access Points, Mobility Switches and Edge Switches) are typically implemented to have a control plane that controls operation of the network element and a data plane that handles traffic flowing through the network. The data plane typically will have a collection of line cards having ports that connect to links on the network. Data is received at a particular port, switched within the data plane, and output at one or more other ports onto other links on the network. To enable the data to be handled quickly, the data plane is typically implemented in hardware so that all of the decisions as to how to handle the data are performed using hardware lookups, etc. The packets are transferred across the network in accordance with a particular protocol, such as the Internet Protocol (IP).

Ports can fail for many reasons, including line card failure, failure of the link connected to the port (e.g. line cut), far-end line card failure, etc. Likewise, the internal forwarding datapath within the network element may fail which may cause a port or set of ports to appear to have failed, or there may be some other failures along the logical/virtual connection to the port’s external peer endpoint. There are numerous reasons why a port may fail.

In the event a port fails, traffic flowing through the port should be diverted to flow out an alternate port to enable connectivity to be restored through the network. To minimize impact on the traffic being handled by the network element, e.g. to minimize downtime and packet loss, the quicker the rerouting of traffic can occur the better. Preferably, it would be advantageous to enable the traffic to fail over to an alternate port in under ten milliseconds (ms). Preferably, the traffic should be spread across the remaining ports rather than all moved from the failing port to a particular designated alternate port to prevent the designated alternate port from being overloaded with traffic.

The current “network overlay” model of integrating Wireless networks with the wired network infrastructure has drawbacks that are alleviated with a newer design called Wireless LAN Split-Plane architecture. In this architecture the basic data forwarding function is pulled out of the Wireless Switches and is incorporated into the existing Routing switches, thus minimizing the monetary cost and network management headaches for Enterprise networks. In such deployments, the Access Point device that performs the Wireless signaling to the Mobile units, is connected via a point-

to-point tunnel to one member of a pair of Routing switches that form a Routed Split Multi-Link Trunking (RSMLT) system. RSMLT clusters are used traditionally in networks in order to provide resiliency, redundancy and fast traffic recovery when links are severed or a switch goes out of service.

Over time, the manner in which network elements handle data has evolved. For example, two or more physical links may extend between a group of network elements and be used collectively as a multi-link trunk (MLT). When the links of an MLT are physically connected to two different network elements, the MLT is referred to as a Split Multi-Link Trunk (SMLT). In particular, each of the links in the MLT may be used by either of the network elements to forward data to the other. Thus, if a first network element has data (e.g., a frame/packet) to send to a second network element, the first network element may select one of the links from the MLT and transmit the packet over that link to the second network element.

As noted above, depending on the manner in which the network elements are interconnected, there may be many ways for the network element to forward a frame/packet to enable the frame/packet to reach its destination. As used herein, the term “cluster” is used to refer to one or more nodes providing node-level resiliency at the network level. Logical connections between the cluster nodes are referred to herein as Inter-Switch Trunks (ISTs). ISTs may be physical links that extend from one network element to a neighboring network element in the cluster, or may be logical links that tunnel through one or more intermediate network elements within the cluster. The node that receives a packet will be referred to as a local node. All other nodes within the cluster are referred to as remote nodes with respect to the received packet.

Two or more links may be grouped to form a Multi-Link Trunk (MLT). Each MLT will be assigned a MLT group ID (MLT-ID). An MLT with all its port members only on the local node is referred to as a local MLT group. An MLT group with some of its port members on the local node and the rest on one or more of the remote nodes is referred to as a Split MLT or SMLT group. An SMLT will be assigned a SMLT ID which is a global value within the cluster and unique across the cluster nodes.

When a logical port is implemented as a MLT or SMLT, there are actually multiple physical ports that are capable of forwarding a packet to its next hop on the network. Accordingly, if one of the ports of a MLT/SMLT fails, it would be advantageous to cause the packet to be forwarded on one of the remaining ports so that the packet can traverse the network rather than being dropped. Likewise, rather than designate a primary and backup port for each port in the SMLT, it would be advantageous to load share the packets across the remaining ports of the MLT/SMLT so that the packets may be distributed across the remaining ports that are UP. According to an embodiment, this process is implemented in hardware so that the fastpath (data plane) can automatically accommodate individual and multiple port failures and automatically redirect packet traffic across the remaining ports in an equitable manner.

Wireless Local Area Network (WLAN) split-plane requires data plane or packet forwarding function of the WLAN end devices (mobility units) traffic being handled by the routing switches in the wired network rather than tunneling all the WLAN traffic to centralized controllers. These switches are called mobility switches and they are capable of terminating tunnels initiated by the WLAN access point (AP) devices. They perform the packet forwarding function for the WLAN network.

A significant number of deployments have (routed) split multi-link trunking (SMLT/RSMLT) in the distribution and

core layers. (R)SMLT provides redundancy and traffic load balancing of the routing switch nodes. Typical WLAN split plane deployment in a customer network will prefer the mobility switch function be performed by a routing switch in the distribution or core layers. This essentially requires inter-working between WLAN split plane and RSMLT.

RSMLT provides resiliency to the network. WLAN split plane is a unique solution to converge the wired and wireless networks. In these converged networks, wired network resiliency capabilities should be extended to the WLAN networks.

SUMMARY

In a world where communications and permanent connectivity of devices to the Internet have become the norm, loss of connectivity and jitter are becoming unacceptable. This requires networks to provide as fast recovery as possible and as small traffic loss as can possibly be allowed. The Virtual Router Redundancy Protocol (VRRP) and its newer incarnation VRRPv3 provide a level of resiliency and recovery from loss which may be acceptable for small networks or applications that are not time critical. VRRPv3 does not address the special issues raised in the case of point to point tunnels like the ones created in a Wireless LAN Split-plane network architecture where additional protocol dependencies come into play.

The current "network overlay" model of integrating Wireless networks with the wired network infrastructure has drawbacks that are alleviated with a newer design called Wireless LAN Split-Plane architecture. In this architecture the basic data forwarding function is pulled out of the Wireless Switches and is incorporated into the existing Routing switches, thus minimizing the monetary cost and network management headaches for Enterprise networks. In such deployments, the Access Point device that performs the Wireless signaling to the Mobile units, is connected via a point-to-point tunnel to one member of a pair of Routing switches that form a Routed Split Multi-Link Trunking (RSMLT) system. RSMLT clusters are used traditionally in networks in order to provide resiliency, redundancy and fast traffic recovery when links are severed or a switch goes out of service.

Conventional mechanisms such as those explained above suffer from a variety of deficiencies. The conventional manner for providing load-sharing services is the VRRP protocol and its different incarnations. In the VRRP case configuration is required in order to perform load sharing. When a link is severed or when the VRRP master switch stops functioning, VRRP on the still active peer depends on timers per interface in order to perform the switch over, which overloads the system as the number of interfaces increases and the timers have short timeouts thus hindering the quick recovery.

It is therefore desirable to extend the RSMLT in order to provide load-sharing and resiliency services to networking topologies that include tunneling. None of the aforementioned cases deals with providing redundancy, load sharing, and quick recovery in the case of WLAN Split-plane Access Tunnels.

Embodiments of the invention significantly overcome such deficiencies. In a WLAN Split-Plane deployment, a RSMLT switch cluster provides redundancy, resiliency and fast traffic recovery when either a link or a switch fails, for data traffic coming from, or destined to, an Access Point wireless switch via an Access Tunnel, regardless of which cluster member is the actual tunnel terminus.

In a particular embodiment of a method for providing routed split multi-link trunking resiliency for wireless local area network split-plane environments, the method includes

learning, by a first network switch, first data associated with a second network switch wherein the first network switch and the second network switch operate concurrently and wherein the first network switch terminates an Access Tunnel (AT). The method further includes learning, by the second network switch, second data associated with the first network switch. Additionally, the method includes synchronizing the first and second data between the first network switch and the second network switch, wherein the data includes MAC Addresses, Internet Protocol (IP) interface addresses, Virtual Local Area Network (VLAN) associations, AT assignments and Access Point IP addresses. The method further includes load sharing tunnel data packets by the first network switch and the second network switch wherein when the second network switch receives tunnel data packets destined to the first network switch the second network acts as a tunnel endpoint and performs tunnel functions for the tunnel data packets, forwarding tunnel control packets received by the second network switch to the first network switch over an Inter Switch Trunk (IST) and detecting a failure relating to the first network switch. Further still the method includes performing recovery actions upon the detecting a failure, wherein the recovery actions include recognizing addresses and tunnels associated with the first network switch as belonging to the second network switch and handling the data packets destined for the first network switch by the second network switch for a predetermined amount of time, and establishing a new AT with the second network switch.

Other embodiments include a computer readable medium having computer readable code thereon for providing routed split multi-link trunking resiliency for wireless local area network split-plane environments. The computer readable medium includes instructions for learning, by a first network switch, first data associated with a second network switch wherein the first network switch and the second network switch operate concurrently and wherein the first network switch terminates an Access Tunnel (AT). The method further includes as learning, by the second network switch, second data associated with the first network switch. Additionally, the computer readable medium includes instructions for synchronizing the first and second data between the first network switch and the second network switch, wherein the data includes MAC Addresses, Internet Protocol (IP) interface addresses, Virtual Local Area Network (VLAN) associations, AT assignments and Access Point IP addresses. The computer readable medium further includes instructions for load sharing tunnel data packets by the first network switch and the second network switch wherein when the second network switch receives tunnel data packets destined to the first network switch the second network acts as a tunnel endpoint and performs tunnel functions for the tunnel data packets, forwarding tunnel control packets received by the second network switch to the first network switch over an Inter Switch Trunk (IST) and detecting a failure relating to the first network switch. Further still the computer readable medium includes instructions for performing recovery actions upon the detecting a failure, wherein the recovery actions include recognizing addresses and tunnels associated with the first network switch as belonging to the second network switch and handling the data packets destined for the first network switch by the second network switch for a predetermined amount of time, and establishing a new AT with the second network switch.

Still other embodiments include a computerized device (e.g., a network switch), configured to process all the method operations disclosed herein as embodiments of the invention. In such embodiments, the computerized device includes a

memory system, a processor, communications interface in an interconnection mechanism connecting these components. The memory system is encoded with a process that provides Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments as explained herein that when performed (e.g. when executing) on the processor, operates as explained herein within the computerized device to perform all of the method embodiments and operations explained herein as embodiments of the invention. Thus any computerized device that performs or is programmed to perform up processing explained herein is an embodiment of the invention.

Other arrangements of embodiments of the invention that are disclosed herein include software programs to perform the method embodiment steps and operations summarized above and disclosed in detail below. More particularly, a computer program product is one embodiment that has a computer-readable medium including computer program logic encoded thereon that when performed in a computerized device provides associated operations providing Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments as explained herein. The computer program logic, when executed on at least one processor with a computing system, causes the processor to perform the operations (e.g., the methods) indicated herein as embodiments of the invention. Such arrangements of the invention are typically provided as software, code and/or other data structures arranged or encoded on a computer readable medium such as an optical medium (e.g., CD-ROM), floppy or hard disk or other a medium such as firmware or microcode in one or more ROM or RAM or PROM chips or as an Application Specific Integrated Circuit (ASIC) or as downloadable software images in one or more modules, shared libraries, etc. The software or firmware or other such configurations can be installed onto a computerized device to cause one or more processors in the computerized device to perform the techniques explained herein as embodiments of the invention. Software processes that operate in a collection of computerized devices, such as in a group of data communications devices or other entities can also provide the system of the invention. The system of the invention can be distributed between many software processes on several data communications devices, or all processes could run on a small set of dedicated computers, or on one computer alone.

It is to be understood that the embodiments of the invention can be embodied strictly as a software program, as software and hardware, or as hardware and/or circuitry alone, such as within a data communications device. The features of the invention, as explained herein, may be employed in data communications devices and/or software systems for such devices such as those manufactured by Avaya, Inc. of Basking Ridge, N.J.

Note that each of the different features, techniques, configurations, etc. discussed in this disclosure can be executed independently or in combination. Accordingly, the present invention can be embodied and viewed in many different ways. Also, note that this summary section herein does not specify every embodiment and/or incrementally novel aspect of the present disclosure or claimed invention. Instead, this summary only provides a preliminary discussion of different embodiments and corresponding points of novelty over conventional techniques. For additional details, elements, and/or possible perspectives (permutations) of the invention, the reader is directed to the Detailed Description section and corresponding figures of the present disclosure as further discussed below.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing will be apparent from the following more particular description of preferred embodiments of the invention, as illustrated in the accompanying drawings in which like reference characters refer to the same parts throughout the different views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention.

FIG. 1 depicts a network for performing Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments in accordance with embodiments of the invention;

FIG. 2 depicts a network for performing Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments with a switch failure;

FIG. 3 depicts a network for performing Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments wherein a new Access tunnel is established;

FIGS. 4A and 4B depict a flow diagram for a particular method for performing Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments in accordance with embodiments of the invention; and

FIG. 5 illustrates an example network switch architecture for a device that performs routed split multi-link trunking resiliency for wireless local area network split-plane in accordance with embodiments of the invention.

DETAILED DESCRIPTION

Referring to FIGS. 1-3, a basic RSMLT network environment 10 is shown. The network environment 10 includes mobile units 28 and 30 in wireless communication with Access Point (AP) 12. Mobile unit 28 is shown as a PDA and mobile unit 30 is shown as a laptop computer although it should be appreciated that other types of mobile units may be used, such as PDAs, cellular devices and the like. AP 12 is in communication with Access Switch 14 which is in communication with network switch 16 and network switch 18. Network switches 16 and 18 are also referred to herein as mobility switches. Access Switch 14, network switch 16 and network switch 18 are connected to each other via a Split Multi-Link Trunk (SMLT) 20 which provides for extra bandwidth as well as redundancy. The network switches 16 and 18 are configured to run Layer 3 protocols (IPv4 family in particular) as well as an additional proprietary protocol via which they exchange information about their configuration, state, databases and link state. This latter protocol runs over the MLT connection between the two peers that form the RSMLT cluster named Inter-Switch Trunk (IST) 22. Network switch 16 and network switch 18 are peers for each other and part of a cluster. On one side the cluster is connected to the network core 26 and on the other side each peer is connected via SMLT 20 to the Access Switch 14 which is running as a Layer 2 switch and is unaware of the fact that it is connecting to two distinct systems. The Access Switch 14 performs the load-sharing function when forwarding traffic over the SMLT links which is flowing towards the network core. The two RSMLT peers in the cluster install each other's MAC addresses, ARP, IPv4 interface addresses, Access Tunnel assignments, Mobility VLAN associations and Access Point information in their own respective tables, thus making themselves capable of routing traffic destined for their peer, that due to the load-sharing function of the access switch could end up on themselves. By routing the traffic themselves as if

the intended recipient, the packets avoid taking an extra hop through the network thus keeping the latency low while providing redundancy.

To provide Mobile units with connectivity to a wired network, the Access Point (AP) performs dual function at the edge of the network. On one side it provides the Radio Frequency (RF) connectivity to the mobile unit and on the other is most commonly wired to a Layer 2 Access Switch **14**. The AP using a standards tracked protocol running over IPv4, connects to a specialized system called the Wireless Controller (WC) (not shown) which carries all the information regarding the switches that serve particular sub-networks (Mobility VLANs). The WC provides the AP **12** with the IPv4 address of the system which will serve as its entrance to the network and this is commonly one of the two peers in the RSMLT cluster. The AP **14** contacts its network server and a tunnel is formed, called Access Tunnel (AT) **24** between the two switches **12** and **16** which flows through the Access Switch **14** and will serve as the conduit for traffic initiated or destined to all the mobile units **28** and **30** served by this Access Point **12**. Control traffic regarding the Access Tunnel **24** is always initiated or terminated by the AP **12** and the RSMLT cluster node **16** that the AP has chosen to be its server, whereas data traffic is actually serviced by both cluster peers **16** and **18**, as if the tunnel **24** had two end-servers. This is achieved by the cluster peers **16** and **18** exchanging the tunnel identification and the AP's IPv4 address. On the RSMLT cluster side, the tunnel space is divided between the two peers **16** and **18** in order to avoid assignment collisions when cluster peers restart after failure, by letting the RSMLT cluster peer with the IP address having higher numerical value when viewed as a single number assign odd numbered tunnels and the one with lower numerically IP address assign even numbered tunnels.

During normal operation, traffic from the mobile unit **28** and or **30** is captured by the Access Point **12** which encapsulates it in a CAPWAP header and sends it via the Access Tunnel **24** to the RSMLT cluster switch **16** that is the tunnel termination. Due to the load sharing functionality of the Access Switch **14** sitting between the two end-points of the tunnel, the traffic is split between the two peers **16** and **18**, but regardless of the original intended recipient, the actual recipient decapsulates the packets and routes them through the network core **26**. This is achieved by having both peers exchange their respective MAC addresses, IPv4 interface addresses and the VLAN associations, Access Tunnel assignments and their respective AP IPv4 address and the fact that each cluster member marks in its hardware its peer's MAC address as an own address. Tunnel related control traffic between the AP **12** and the tunnel terminal member of the cluster **16**, is always forwarded to and processed by the actual cluster member **16** that is the tunnel owner even if the traffic takes an extra hop via its cluster peer **18** due to the load sharing actions of Access Switch **14**.

Referring now to FIG. 2, in the event that the RSMLT cluster switch **16** that owns the Access Tunnel **24** fails at some point or is taken off service for a period of time, the active cluster peer **18** keeps forwarding the data traffic on its behalf. This is accomplished by network switch **18** continuing for a time to assume the identity of the inactive switch **16** and recognizing the inactive switch's addresses and tunnel as its own. The switchover in the case of failure is as immediate as the Access Switch **14** can recognize the fact that its connections to that member of the RSMLT cluster are severed and redirects all incoming traffic to the active member **18**, which usually is in the order of milliseconds. Data traffic thus continues to flow unabated in either direction.

Referring now in particular to FIG. 3, when the AP **12** realizes that the tunnel's cluster owner **16** is not responding to the tunnel control packets, the AP **12** establishes a new access tunnel **32** with the still active cluster member **18** and redirects all traffic via this new tunnel **32**. In a similar fashion the active cluster member **18** switches over the traffic intended for the Access Point via the newly formed tunnel **32**.

The presently described method for performing Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments provides three basic services. These services are redundancy of links and switches (nodes), load-sharing, and fast recovery in case of loss of link or switch for WLAN traffic over a tunnel to and from an Access Point wireless device.

Load sharing comes for free since it is actually performed by the Layer 2 switch connecting to the two RSMLT cluster switches via a MLT. In the RSMLT case (as opposed to the earlier described VRRP technique) the switch over is almost instant due to the fact that there is no timer dependency and the two peers in the cluster always forward routed traffic regardless of which one was the intended gateway.

FIGS. 4A and 4B depict a flow diagram for a particular method for performing Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments in accordance with embodiments of the invention. The rectangular elements are herein denoted "processing blocks" and represent computer software instructions or groups of instructions. Alternatively, the processing blocks represent steps performed by functionally equivalent circuits such as a digital signal processor circuit or an application specific integrated circuit (ASIC). The flow diagrams do not depict the syntax of any particular programming language. Rather, the flow diagrams illustrate the functional information one of ordinary skill in the art requires to fabricate circuits or to generate computer software to perform the processing required in accordance with the present invention. It should be noted that many routine program elements, such as initialization of loops and variables and the use of temporary variables are not shown. It will be appreciated by those of ordinary skill in the art that unless otherwise indicated herein, the particular sequence of steps described is illustrative only and can be varied without departing from the spirit of the invention. Thus, unless otherwise stated the steps described below are unordered meaning that, when possible, the steps can be performed in any convenient or desirable order.

Referring now to FIGS. 4A and 4B, a particular embodiment of a method **100** for performing Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments is shown. Method **100** begins with processing block **102** which discloses learning, by a first network switch, first data associated with a second network switch, wherein the first network switch and the second network switch operate concurrently and wherein the first network switch terminates an Access Tunnel (AT).

Processing block **104** recites learning, by the second network switch, second data associated with the first network switch.

Processing block **106** discloses synchronizing the first and second data between the first network switch and the second network switch, wherein the data includes MAC Addresses, Internet Protocol (IP) interface addresses, Virtual Local Area Network (VLAN) associations, AT assignments and Access Point IP addresses.

Processing block **108** states load sharing tunnel data packets by the first network switch and the second network switch wherein when the second network switch receives tunnel data

packets destined to the first network switch the second network acts as a tunnel endpoint and performs tunnel functions for the tunnel data packets.

Processing block 110 recites forwarding tunnel control packets received by the second network switch to the first network switch over the IST.

Processing continues with processing block 112 which discloses detecting a failure relating to the first network switch. As shown in processing block 14, detecting a failure relating to the first network switch comprises detecting one of the group consisting of a link failure and a switch failure. Alternately, as shown in processing block 116 detecting a failure relating to the first network switch comprises determining that the first network switch is not responding to tunnel control packets.

Processing block 118 states performing recovery actions upon the detecting a failure. As shown in processing block 120, the recovery actions include recognizing addresses and tunnels associated with the first network switch as belonging to the second network switch and handling the data packets destined for the first network switch by the second network switch for a predetermined amount of time. As further shown in processing block 122, the recovery actions further include establishing a new AT with the second network switch, and, as shown in processing block 124, forwarding packets destined for the first network switch via the new AT to the second network switch.

FIG. 5 is a block diagram illustrating example architecture of a computer system (e.g. a network switch) 210 that executes, runs, interprets, operates or otherwise performs Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments (RSMLT for WLAN) application 240-1 and Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments process 240-2 suitable for use in explaining example configurations disclosed herein. As shown in this example, the computer system 210 includes an interconnection mechanism 211 such as a data bus or other circuitry that couples a memory system 212, a processor 213, an input/output interface 214, and a communications interface 215. The communications interface 215 enables the computer system 210 to communicate with other devices (i.e., other computers) on a network (not shown).

The memory system 212 is any type of computer readable medium, and in this example, is encoded with a Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments application 240-1 as explained herein. The Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments application 240-1 may be embodied as software code such as data and/or logic instructions (e.g., code stored in the memory or on another computer readable medium such as a removable disk) that supports processing functionality according to different embodiments described herein. During operation of the computer system 210, the processor 213 accesses the memory system 212 via the interconnect 211 in order to launch, run, execute, interpret or otherwise perform the logic instructions of a Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments application 240-1. Execution of a Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments application 240-1 in this manner produces processing functionality in the Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments process 240-2. In other words, the Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments process 240-2

represents one or more portions or runtime instances of a Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments application 240-1 (or the entire a Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments application 240-1) performing or executing within or upon the processor 213 in the computerized device 210 at runtime.

It is noted that example configurations disclosed herein include the Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments application 240-1 itself (i.e., in the form of un-executed or non-performing logic instructions and/or data). The Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments application 240-1 may be stored on a computer readable medium (such as a floppy disk), hard disk, electronic, magnetic, optical, or other computer readable medium. A Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments application 240-1 may also be stored in a memory system 212 such as in firmware, read only memory (ROM), or, as in this example, as executable code in, for example, Random Access Memory (RAM). In addition to these embodiments, it should also be noted that other embodiments herein include the execution of a Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments application 240-1 in the processor 213 as the Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments process 240-2. Those skilled in the art will understand that the computer system 210 may include other processes and/or software and hardware components, such as an operating system not shown in this example.

During operation, processor 213 of computer system 200 accesses memory system 212 via the interconnect 211 in order to launch, run, execute, interpret or otherwise perform the logic instructions of the Routed split multi-link trunking resiliency for wireless local area network split-plane environments application 240-1. Execution of Routed split multi-link trunking resiliency for wireless local area network split-plane environments application 240-1 produces processing functionality in Routed split multi-link trunking resiliency for wireless local area network split-plane environments process 240-2. In other words, the Routed split multi-link trunking resiliency for wireless local area network split-plane environments process 240-2 represents one or more portions of the Routed split multi-link trunking resiliency for wireless local area network split-plane environments application 240-1 (or the entire application) performing within or upon the processor 213 in the computer system 200.

It should be noted that, in addition to the Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments process 240-2, embodiments herein include the Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments application 240-1 itself (i.e., the un-executed or non-performing logic instructions and/or data). The Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments application 240-1 can be stored on a computer readable medium such as a floppy disk, hard disk, or optical medium. The Routed Split Multi-Link Trunking Resiliency for Wireless Local Area Network split-plane environments application 240-1 can also be stored in a memory type system such as in firmware, read only memory (ROM), or, as in this example, as executable code within the memory system 212 (e.g., within Random Access Memory or RAM).

In addition to these embodiments, it should also be noted that other embodiments herein include the execution of Routed split multi-link trunking resiliency for wireless local area network split-plane environments application **240-1** in processor **213** as the Routed split multi-link trunking resiliency for wireless local area network split-plane environments process **240-2**. Those skilled in the art will understand that the computer system **200** can include other processes and/or software and hardware components, such as an operating system that controls allocation and use of hardware resources associated with the computer system **200**.

References to “a microprocessor” and “a processor”, or “the microprocessor” and “the processor,” may be understood to include one or more microprocessors that may communicate in a stand-alone and/or a distributed environment(s), and may thus be configured to communicate via wired or wireless communications with other processors, where such one or more processor may be configured to operate on one or more processor-controlled devices that may be similar or different devices. Use of such “microprocessor” or “processor” terminology may thus also be understood to include a central processing unit, an arithmetic logic unit, an application-specific integrated circuit (IC), and/or a task engine, with such examples provided for illustration and not limitation.

Furthermore, references to memory, unless otherwise specified, may include one or more processor-readable and accessible memory elements and/or components that may be internal to the processor-controlled device, external to the processor-controlled device, and/or may be accessed via a wired or wireless network using a variety of communications protocols, and unless otherwise specified, may be arranged to include a combination of external and internal memory devices, where such memory may be contiguous and/or partitioned based on the application. Accordingly, references to a database may be understood to include one or more memory associations, where such references may include commercially available database products (e.g., SQL, Informix, Oracle) and also proprietary databases, and may also include other structures for associating memory such as links, queues, graphs, trees, with such structures provided for illustration and not limitation.

References to a network, unless provided otherwise, may include one or more intranets and/or the Internet, as well as a virtual network. References herein to microprocessor instructions or microprocessor-executable instructions, in accordance with the above, may be understood to include program-mable hardware.

Unless otherwise stated, use of the word “substantially” may be construed to include a precise relationship, condition, arrangement, orientation, and/or other characteristic, and deviations thereof as understood by one of ordinary skill in the art, to the extent that such deviations do not materially affect the disclosed methods and systems.

Throughout the entirety of the present disclosure, use of the articles “a” or “an” to modify a noun may be understood to be used for convenience and to include one, or more than one of the modified noun, unless otherwise specifically stated.

Elements, components, modules, and/or parts thereof that are described and/or otherwise portrayed through the figures to communicate with, be associated with, and/or be based on, something else, may be understood to so communicate, be associated with, and/or be based on in a direct and/or indirect manner, unless otherwise stipulated herein.

Although the methods and systems have been described relative to a specific embodiment thereof, they are not so limited. Obviously many modifications and variations may become apparent in light of the above teachings. Many addi-

tional changes in the details, materials, and arrangement of parts, herein described and illustrated, may be made by those skilled in the art.

Having described preferred embodiments of the invention it will now become apparent to those of ordinary skill in the art that other embodiments incorporating these concepts may be used. Additionally, the software included as part of the invention may be embodied in a computer program product that includes a computer useable medium. For example, such a computer useable medium can include a readable memory device, such as a hard drive device, a CD-ROM, a DVD-ROM, or a computer diskette, having computer readable program code segments stored thereon. The computer readable medium can also include a communications link, either optical, wired, or wireless, having program code segments carried thereon as digital or analog signals. Accordingly, it is submitted that that the invention should not be limited to the described embodiments but rather should be limited only by the spirit and scope of the appended claims.

What is claimed is:

1. A method comprising:

learning, by a first network switch, first data associated with a second network switch wherein said first network switch and said second network switch operate concurrently and wherein said first network switch terminates an Access Tunnel (AT);

learning, by said second network switch, second data associated with said first network switch;

synchronizing said first and second data between said first network switch and said second network switch, wherein said first and second data includes MAC Addresses, Internet Protocol (IP) interface addresses, Virtual Local Area Network (VLAN) associations, AT assignments and Access Point IP addresses; and

load sharing tunnel data packets by said first network switch and said second network switch wherein when said second network switch receives tunnel data packets destined to said first network switch said second network switch acts as a tunnel endpoint and performs tunnel functions for said tunnel data packets;

forwarding tunnel control packets received by said second network switch to said first network switch over an Inter Switch Trunk (IST);

detecting a failure relating to said first network switch, wherein said detecting a failure relating to said first network switch further comprises performing recovery actions upon said detecting a failure,

wherein said recovery actions include recognizing addresses and tunnels associated with said first network switch as belonging to said second network switch and handling said data packets destined for said first network switch by said second network switch for a predetermined amount of time; and

establishing a new AT with said second network switch.

2. The method of claim 1 further comprising forwarding packets destined for said first network switch via said new AT to said second network switch.

3. The method of claim 1 wherein said detecting a failure relating to said first network switch comprises detecting one of the group consisting of a link failure and a switch failure.

4. The method of claim 1 wherein said detecting a failure relating to said first network switch comprises determining that the first network switch is not responding to tunnel control packets.

5. A first network switch comprising:

a memory; a processor; a communications interface;

13

an interconnection mechanism coupling the memory, the processor and the communications interface; and wherein the memory is encoded with an application providing routed split multi-link trunking resiliency for wireless local area network split-plane environments, that when performed on the processor, provides a process for processing information, the process causing the first network switch to perform the operations of:

learning data associated with a second network switch, wherein said data includes MAC Addresses, Internet Protocol (IP) interface addresses, Virtual Local Area Network (VLAN) associations, AT assignments and Access Point IP addresses,

wherein said first network switch and said second network switch operate concurrently and wherein said second network switch terminates an Access Tunnel (AT);

load sharing tunnel data packets,

wherein when said first network switch receives tunnel data packets destined to said second network switch, acting as a tunnel endpoint and performing tunnel functions for said tunnel data packets;

forwarding tunnel control packets received by said first network switch to said second network switch over an Inter Switch Trunk (IST);

becoming aware of a failure relating to said second network switch,

wherein said becoming aware of a failure relating to said second network switch further comprises performing recovery actions,

wherein said recovery actions include recognizing addresses and tunnels associated with said second network switch as belonging to said first network switch and handling said data packets destined for said second network switch by said first network switch for a predetermined amount of time; and

becoming part of a new AT.

6. The first network switch of claim 5 further comprising receiving and processing packets destined for said second network switch via said new AT.

7. The first network switch of claim 5 wherein said becoming aware of a failure relating to said second network switch comprises detecting one of the group consisting of a link failure and a switch failure.

8. The first network switch of claim 5 wherein said becoming aware of a failure relating to said second network switch comprises determining that the second network switch is not responding to tunnel control packets.

9. A non-transitory computer readable storage medium having computer readable code thereon for providing routed split multi-link trunking resiliency for wireless local area

14

network split-plane environment, the medium including instructions in which a computer system performs operations comprising:

learning, by a first network switch, first data associated with a second network switch wherein said first network switch and said second network switch operate concurrently and wherein said first network switch terminates an Access Tunnel (AT);

learning, by said second network switch, second data associated with said first network switch;

synchronizing said first and second data between said first network switch and said second network switch, wherein said first and second data includes MAC Addresses, Internet Protocol (IP) interface addresses, Virtual Local Area Network (VLAN) associations, AT assignments and Access Point IP addresses; and

load sharing tunnel data packets by said first network switch and said second network switch wherein when said second network switch receives tunnel data packets destined to said first network switch said second network switch acts as a tunnel endpoint and performs tunnel functions for said tunnel data packets;

forwarding tunnel control packets received by said second network switch to said first network switch over an Inter Switch Trunk (IST);

detecting a failure relating to said first network switch,

wherein said detecting a failure relating to said first network switch further comprises performing recovery actions upon said detecting a failure,

wherein said recovery actions include recognizing addresses and tunnels associated with said first network switch as belonging to said second network switch and handling said data packets destined for said first network switch by said second network switch for a predetermined amount of time; and

establishing a new AT with said second network switch.

10. The non-transitory computer readable storage medium of claim 9 further comprising forwarding packets destined for said first network switch via said new AT to said second network switch.

11. The non-transitory computer readable storage medium of claim 9, wherein said detecting a failure relating to said first network switch comprises detecting one of the group consisting of a link failure, a switch failure, and determining that the first network switch is not responding to tunnel control packets.

* * * * *